



[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search: ☒ The ACM Digital Library ☐ The Guide

+backup, +archive, +data +synchronization, +data +duplication



THE ACM DIGITAL LIBRARY

[Feedback](#) [Report a problem](#) [Satisfaction s](#)

Terms used

[backup](#) [archive](#) [data](#) [synchronization](#) [data](#) [duplication](#) [backup](#) [backing up](#) [archiving](#) [archived](#) [synchronizati](#)

Sort results by

[Save results to a Binder](#)

[Try an Advanced Search](#)

Display results

[Search Tips](#)

[Try this search in The ACM Gui](#)

☐ [Open results in a new window](#)

Results 1 - 20 of 75

Result page: [1](#) [2](#) [3](#) [4](#) [next](#)

Relevance scale

- 1 [Industrial sessions: beyond relational tables: Coordinating backup/recovery and data consistency between database and file systems](#)



Suparna Bhattacharya, C. Mohan, Karen W. Brannon, Inderpal Narang, Hui-I Hsiao, Mahadevan Subramanian

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02**

Publisher: ACM Press

Full text available: [pdf\(1.44 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Managing a combined store consisting of database data and file data in a robust and consistent manner is a challenge for database systems and content management systems. In such a hybrid system, images, videos, engineering drawings, etc. are stored as files on a file server while metadata referencing/indexing such files is created and stored in a relational database to take advantage of efficient search. In this paper we describe solutions for two potentially problematic aspects of such data ...

Keywords: DB2, content management, database backup, database recovery, datalinks

- 2 [The LHAM log-structured history data access method](#)

Peter Muth, Patrick O'Neil, Achim Pick, Gerhard Weikum

February 2000 **The VLDB Journal — The International Journal on Very Large Data Bases**, Vol. 9, Issue 3-4

Publisher: Springer-Verlag New York, Inc.

Full text available: [pdf\(494.76 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

Numerous applications such as stock market or medical information systems require that both historical and current data be logically integrated into a temporal database. The underlying access method must support different forms of "time-travel" queries, the migration of old record versions onto inexpensive archive media, and high insertion and update rates. This paper presents an access method for transaction-time temporal data, called the log-structured history data access method.

Keywords: Data warehouses, Index structures, Performance, Storage systems, Temporal data

- 3 [Personal information management: To have and to hold: exploring the personal archive](#)



Joseph 'Jofish' Kaye, Janet Vertesi, Shari Avery, Allan Dafoe, Shay David, Lisa Onaga, Ivan Rosero, Trevor Pinch

April 2006 **Proceedings of the SIGCHI conference on Human Factors in computing system '06**

Publisher: ACM Press

Full text available:  [pdf\(761.50 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The personal archive is not only about efficient storage and retrieval of information. This paper describes a study of forty-eight academics and the techniques and tools they use to manage the digital and material archiving of papers, emails, documents, internet bookmarks, correspondence, other artifacts. We present two sets of results: we first discuss rationales behind subjects' archiving which go beyond information retrieval to include creating a legacy, sharing resources, confronting

Keywords: archiving, bookmarks, email, ethnography, filing, identity

4 Peer-to-peer infrastructure: Pastiche: making backup cheap and easy



Landon P. Cox, Christopher D. Murray, Brian D. Noble

December 2002 **ACM SIGOPS Operating Systems Review**, Volume 36 Issue SI

Publisher: ACM Press

Full text available:  [pdf\(1.65 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#)

Backup is cumbersome and expensive. Individual users almost never back up their data, and backup has a significant cost in large organizations. This paper presents *Pastiche*, a simple and inexpensive backup system. Pastiche exploits excess disk capacity to perform peer-to-peer backup with no administrative costs. Each node minimizes storage overhead by selecting peers that share a significant amount of data. It is easy for common installations to find suitable peers, and peers with high overlap

5 Principles of transaction-oriented database recovery



Theo Haerder, Andreas Reuter

December 1983 **ACM Computing Surveys (CSUR)**, Volume 15 Issue 4

Publisher: ACM Press

Full text available:  [pdf\(2.48 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

6 ARIES: a transaction recovery method supporting fine-granularity locking and partial rollback using write-ahead logging



C. Mohan, Don Haderle, Bruce Lindsay, Hamid Pirahesh, Peter Schwarz

March 1992 **ACM Transactions on Database Systems (TODS)**, Volume 17 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(5.23 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

DB2TM, IMS, and TandemTM systems. ARIES is applicable not only to database management systems but also to persistent object-oriented languages, recoverable file systems and transaction-based operating systems. ARIES has been implemented, to varying degrees, in IBM's OS/2TM Extended Edition Database Manager, DB2, Workstation Data Save Facility/VM, Starburst and QuickSilver, and the University of Wisconsin's EXODUS and Gamma d ...

Keywords: buffer management, latching, locking, space management, write-ahead logging

7 File servers for network-based distributed systems



Liba Svobodova

December 1984 **ACM Computing Surveys (CSUR)**, Volume 16 Issue 4

Publisher: ACM Press

Full text available:  [pdf\(4.23 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

8 Data base directions: the next steps



John L. Berg

November 1976 **ACM SIGMOD Record , ACM SIGMIS Database**, Volume 8 , 8 Issue 4 , 2

Publisher: ACM Press

Full text available: pdf(9.95 MB)

Additional Information: [full citation](#), [abstract](#)

What information about data base technology does a manager need to make prudent decisions using this new technology? To provide this information the National Bureau of Standards and the Association for Computing Machinery established a workshop of approximately 80 experts in five subject areas. The five subject areas were auditing, evolving technology, government regulation standards, and user experience. Each area prepared a report contained in these proceedings. The proceedings p ...

Keywords: DBMS, auditing, cost/benefit analysis, data base, data base management, government regulation, management objectives, privacy, security, standards, technology assessment, user experience

9 Disaster recovery techniques for database systems



Manhoi Choy, Hong Va Leong, Man Hon Wong

November 2000 **Communications of the ACM**

Publisher: ACM Press

Full text available: pdf(412.04 KB)

Additional Information: [full citation](#), [references](#), [index terms](#)

10 Illustrative risks to the public in the use of computer systems and related technology



Peter G. Neumann

January 1996 **ACM SIGSOFT Software Engineering Notes**, Volume 21 Issue 1

Publisher: ACM Press

Full text available: pdf(2.54 MB)

Additional Information: [full citation](#)

11 A database cache for high performance and fast restart in database systems



Klaus Elhardt, Rudolf Bayer

December 1984 **ACM Transactions on Database Systems (TODS)**, Volume 9 Issue 4

Publisher: ACM Press

Full text available: pdf(1.72 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Performance in database systems is strongly influenced by buffer management and transaction recovery methods. This paper presents the principles of the database cache, which replaces the traditional buffer. In comparison to buffer management, cache management is more carefully coordinated with transaction management, and integrates transaction recovery. High throughput small- and medium-sized transactions is achieved by fast commit processing and low database t Very fas ...

12 Recovery Techniques for Database Systems



Joost S. M. Verhofstad


June 1978 **ACM Computing Surveys (CSUR)**, Volume 10 Issue 2

Publisher: ACM Press

Full text available: pdf(2.32 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

13 A quantitative analysis of cache policies for scalable network file systems

 Michael D. Dahlin, Clifford J. Mather, Randolph Y. Wang, Thomas E. Anderson, David A. Patterson
May 1994 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1994 ACM SIGMETRICS conference on Measurement and modeling of computer systems**
SIGMETRICS '94, Volume 22 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(1.42 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Current network file system protocols rely heavily on a central server to coordinate file activity among client workstations. This central server can become a bottleneck that limits scalability for environments with large numbers of clients. In central server systems such as NFS and AFS, all client writes, cache misses, and coherence messages are handled by the server. To keep up with this workload, expensive server machines are needed, configured with high-performance CPUs, memory systems, ...

14 Query evaluation techniques for large databases

 Goetz Graefe
June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

Publisher: ACM Press


Full text available:  [pdf\(9.37 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)



Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate sets of complex objects as efficiently as today's database systems manipulate simple records, query processing ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

15 ReVive: cost-effective architectural support for rollback recovery in shared-memory multiprocessors

 Milos Prvulovic, Zheng Zhang, Josep Torrellas
May 2002 **ACM SIGARCH Computer Architecture News , Proceedings of the 29th annual international symposium on Computer architecture ISCA '02 , Proceedings of the 29th annual international symposium on Computer architecture ISCA '02**, Volume 32 Issue 2

Publisher: IEEE Computer Society, ACM Press

Full text available:  [pdf\(1.38 MB\)](#) 
[Publisher Site](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents ReVive, a novel general-purpose rollback recovery mechanism for shared-memory multiprocessors. ReVive carefully balances the conflicting requirements of availability, performance, and hardware cost. ReVive performs checkpointing, logging, and distributed parity protection, all memory-based. It enables recovery from a wide class of errors, including the permanent loss of an entire node. To maintain high performance, ReVive includes specialized hardware that performs frequent operations ...

Keywords: fault tolerance, shared-memory multiprocessors, rollback recovery, recovery, BER, logging, parity, checkpointing, availability

- 16 RPC-V: Toward Fault-Tolerant RPC for Internet Connected Desktop Grids with Volatile No
Samir Djilali, Thomas Herault, Oleg Lodygensky, Tangui Morlier, Gilles Fedak, Franck Cappello
November 2004 **Proceedings of the 2004 ACM/IEEE conference on Supercomputing**

Publisher: IEEE Computer Society

Full text available:  pdf(230.75 KB)

Additional Information: [full citation](#), [abstract](#)

RPC is one of the programming models envisioned for the Grid. In Internet connected Large Scale Grids such as Desktop Grids, nodes and networks failures are not rare events. This paper provides several contributions, examining the feasibility and limits of fault-tolerant RPC on these platforms. First, we characterize these Grids from their fundamental features and demonstrate that their applications scope should be safely restricted to stateless services. Second, we present a new fault-tolerant ...

- 17 Recovery management in QuickSilver



Rober Haskin, Yoni Malachi, Gregory Chan

February 1988 **ACM Transactions on Computer Systems (TOCS)**, Volume 6 Issue 1

Publisher: ACM Press

Full text available:  pdf(2.21 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

This paper describes QuickSilver, developed at the IBM Almaden Research Center, which uses atomic transactions as a unified failure recovery mechanism for a client-server structured distributed system. Transactions allow failure atomicity for related activities at a single server or at a number of independent servers. Rather than bundling transaction management into a dedicated language or recoverable object manager, Quicksilver exposes the basic commit protocol and log recovery ...

- 18 Experience Using Multiprocessor Systems—A Status Report



Anita K. Jones, Peter Schwarz

June 1980 **ACM Computing Surveys (CSUR)**, Volume 12 Issue 2

Publisher: ACM Press

Full text available:  pdf(4.48 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

- 19 Compiler transformations for high-performance computing



David F. Bacon, Susan L. Graham, Oliver J. Sharp

December 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 4

Publisher: ACM Press

Full text available:  pdf(6.32 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

In the last three decades a large number of compiler transformations for optimizing programs have been implemented. Most optimizations for uniprocessors reduce the number of instructions executed by the program using transformations based on the analysis of scalar quantities and data-flow techniques. In contrast, optimizations for high-performance superscalar, vector, and parallel processors maximize parallelism and memory locality with transformations that rely on tracking properties of ...

Keywords: compilation, dependence analysis, locality, multiprocessors, optimization, parallelism, superscalar processors, vectorization

- 20 Illustrative risks to the public in the use of computer systems and related technology



Peter G. Neumann

January 1992 **ACM SIGSOFT Software Engineering Notes**, Volume 17 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(1.65 MB\)](#)

Additional Information: [full citation](#), [citings](#), [index terms](#)

Results 1 - 20 of 75

Result page: [1](#) [2](#) [3](#) [4](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	24454802	@AD<"20030801"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:28
L2	135	(Wai near2 Lam).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:28
L3	9	(Ronald near2 Niles).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:29
L4	0	(Xiaowel near2 Li).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:30
L5	9	(Xiaowei near2 Li).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:31
L6	0	1 and 2 and 3 and 5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:38
L7	24454821	1 or 2 or 3 or 5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:32
L8	267096	(back\$6 near "up") or backup\$4 or archiv\$7 or backstor\$4 or (reserv\$5 adj stor\$4)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:33
L9	9002	data adj synchroniz\$4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:34

EAST Search History

L10	16139	data adj synchroniz\$8	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:34
L11	1	((first and second) adj2 digest\$4) near3 (random adj value)) same (encod\$4 or decod\$4)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:40
L12	201564	(duplicated or duplicate or duplication or duplicating or duplicatted)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:37
L13	18364	(duplicated or duplicate or duplication or duplicating or duplicatted) same (compar\$5)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:37
L14	20333	(primary or secondary) near2 storage	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:38
L15	1	2 and 3 and 5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:38
L16	144	2 or 3 or 5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:39
L17	239	13 and 14	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:39
L18	109	17 and 8	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:39
L19	2	18 and 10	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:39

EAST Search History

L20	921973	(encod\$4 or decod\$4)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:40
L21	62	18 and 20	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:40
L22	1	((first and second) adj2 digest\$4) near3 (random adj value))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:40
L23	3820	(random adj value)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:40
L24	3	21 and 23	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:41
L25	1	22 and 24	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:41
L26	8	1 and 3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:41
L27	2	1 and 24	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/07/18 15:41

[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alt](#)

Welcome United States Patent and Trademark Office

☐ Search Session History[BROWSE](#)[SEARCH](#)[IEEE XPLORE GUIDE](#)

Tue, 18 Jul 2006, 7:45:20 PM EST

Edit an existing query or compose a new query in the Search Query Display.

Search Query Display

Select a search number (#) to:

- Add a query to the Search Query Display
- Combine search queries using AND, OR, or NOT
- Delete a search
- Run a search

Recent Search Queries

- | | |
|-----|--|
| #1 | (((data synchronization) and backup)<in>metadata) |
| #2 | (((data synchronization) and backup)<in>metadata) |
| #3 | ((backup or (backing up) or (backed up) or archiv*)<IN>metadata) |
| #4 | ((first or second) and digest*<IN>metadata) |
| #5 | (random and (value or number)<IN>metadata) |
| #6 | (data and (encod* or decod*)<IN>metadata) |
| #7 | (data and (duplicated or duplicate or duplication or duplicatted or duplicating)<IN>metadata) |
| #8 | (primary and storage<IN>metadata) |
| #9 | (secondary and storage<in>metadata) |
| #10 | (((data synchronization) and backup)<in>metadata) <AND> (((backup or (backing up) or (backed up) or archiv*)<IN>metadata)) |
| #11 | (((first or second) and digest*<IN>metadata) <AND> ((random and (value or number)<IN>metadata)) |
| #12 | ((data and (encod* or decod*)<IN>metadata) <AND> ((data and (duplicated or duplicate or duplication or duplicatted or duplicating)<IN>metadata)) |
| #13 | (((first or second) and digest*<IN>metadata) <AND> ((random and (value or number)<IN>metadata)) <AND> (((data and (encod* or decod*)<IN>metadata) <AND> ((data and (duplicated or duplicate or duplication or duplicatted or duplicating)<IN>metadata))) |

Indexed by
 Inspec

[Help](#) [Contact Us](#) [Privacy](#)

© Copyright 2006 IE